HighSpeed TCP: Some Results and Observations

Evandro de Souza
Distributed System Department
Lawrence Berkeley National Laboratory

May 2003

Collaborators: Deb Agarwal and Sally Floyd

1. Summary

- What is it?
- Topics of investigation
- Methodology
- Experiments
- Issues for the deployment of HighSpeed TCP
- Conclusion

2. What is it?

- HighSpeed TCP for Large Congestion Windows is a modification of TCP's congestion control mechanism to improve performance of TCP connections with large congestion windows.
- HighSpeed TCP is designed to have a different response in environments of very low congestion event rate, and to have the standard TCP response in environments with packet loss rates of at most 10^{-3}

Congestion Avoidance

$$\textbf{ACK} \quad : \quad CWND \leftarrow CWND + \frac{a(CWND)}{CWND} \tag{1}$$

$$\textbf{DROP} \quad : \quad CWND \leftarrow CWND - b(CWND) \times CWND \tag{2}$$

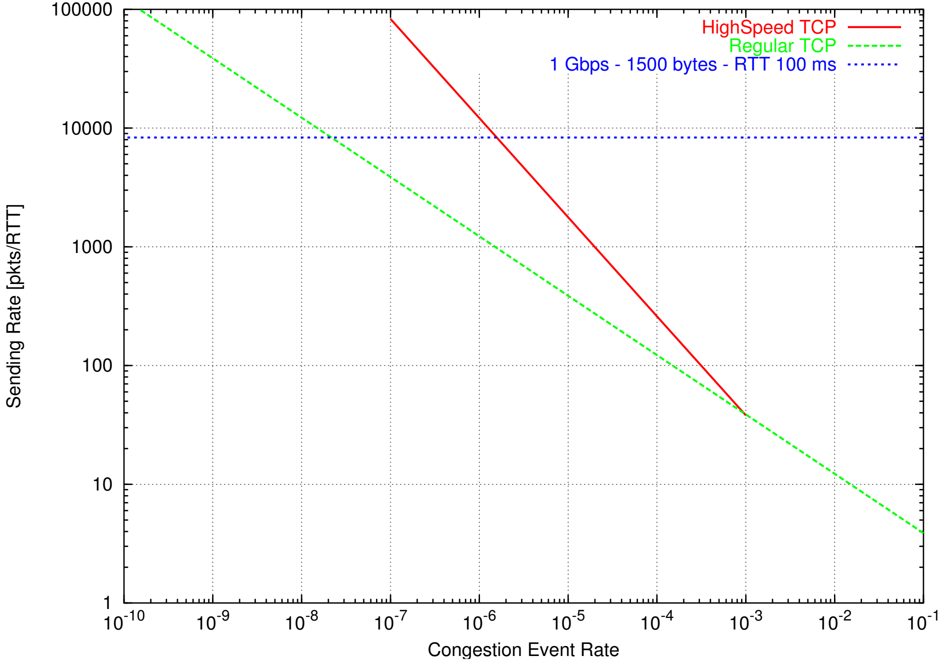

Figure 1: HighSpeed TCP Response Function

3. Topics of Investigation

- What is the behavior of HighSpeed TCP in situations where Regular TCP underperforms
- Is it possible to use HighSpeed TCP together with Regular TCP and maintain an acceptable fairness
- What is the effect of the router queuing policy (RED and DT) on the performance of HighSpeed TCP and on the fairness between HighSpeed TCP and Regular TCP
- Can HighSpeed TCP be a substitute to other types of bulk data transfer (parallel streams)

4. Methodology

4.1. Environment

- single bottleneck
- 1 Gbps / 50 ms
- routers with RED/ECN or DT

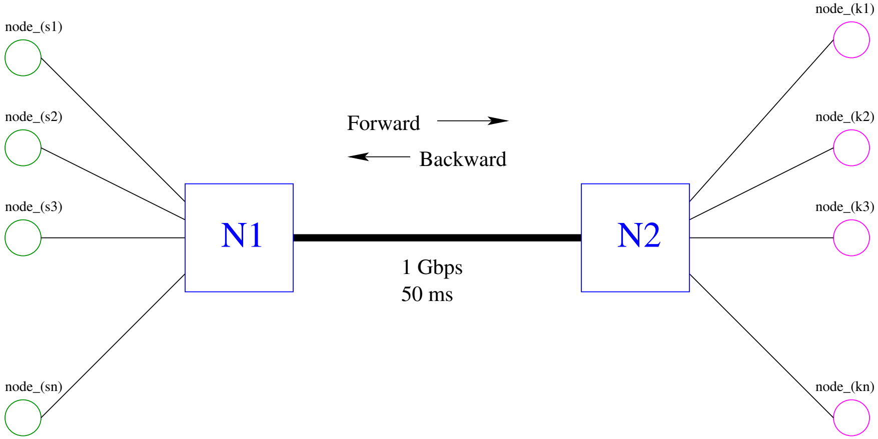

Figure 2: Network Topology

4.2. TCP Flows

- packet size 1500 bytes
- window size large enough to not impose limits
- random times between each packet sent (phase effects)
- Limited Slow-Start (avoid large losses)
- SACK
- FTP long-lived flows
- forward direction

4.3. Background Noise

- web-like flows
- small TCP flows

4.4. Data Collection

- each experiment run 10 times
- 300 simulation seconds each time
- line shown is the median
- around 4900 simulations
- each simulation lasted from 1 hour to 4 days

4.5. Sets of Flows

- only HSTCP flows
- only REGTCP flows
- mix of HSTCP and REGTCP flows

4.6. Network Conditions

- **Ideal Condition:** no other traffic sources and no extra interference beyond that generated by the REGTCP and HSTCP flows
- **Lossy Link Condition:** some number of packets were randomly dropped from the flows, with a defined drop rate
- **Bursty Traffic Condition:** perturbation in the form of bursty flows were randomly distributed throughout the simulation time

5. Experiments

Isolated Flows

- a single REGTCP and a single HSTCP flow run in isolation

Ideal Condition

- HSTCP and REGTCP sets of 1, 2, 6, 10, 20, 30 and 40 flows, and set with a mix of them

Lossy Link Condition

- first set: 10 HSTCP flows, second set: 10 REGTCP flows and third set: mix of 5 HSTCP and 5 REGTCP flows
- link loss rate 0, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}

Bursty Traffic Condition

- first set: 10 HSTCP flows, second set: 10 REGTCP flows and third set: mix of 5 HSTCP and 5 REGTCP flows
- 0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 60, 70 perturbations flows

Competition among Heterogeneous Flows

- 1 HSTCP flow competing against 1, 3, 5, 7, 11, 15 and 19 REGTCP flows
- no external interference

Constant Link Loss of 10^{-5}

- HSTCP and REGTCP sets of 1, 2, 6, 10, 20, 30 and 40 flows, and set with a mix of them
- fixed link loss rate of 10^{-5}

Parallel Streams on Lossy Link Condition

- two sets of flows. The first contained 10 long-lived REGTCP flows and 1, 4, 7, 10, 20 or 30 parallel REGTCP streams. The second formed by the same 10 long-lived REGTCP flows and one HSTCP flow
- link loss rate 0, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}

Parallel Streams on Bursty Traffic Condition

- two sets of flows. The first contained 10 long-lived REGTCP flows and 1, 4, 7, 10, 20 or 30 parallel REGTCP streams. The second formed by the same 10 long-lived REGTCP flows and one HSTCP flow
- 0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 60, 70 perturbations flows

5.1. Isolated Flows

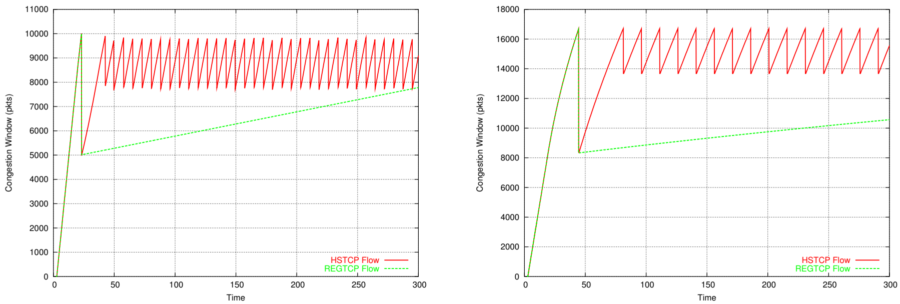

(a) RED (b) DT

Figure 3: **Evolution of Congestion Window for a Single Flow**

- REGTCP flow has a slower growth compared with HSTCP flow
- HSTCP has an oscillatory behavior with a very short period
- strong influence of router queue management in both types of flow

Summary
What is it?
Topics of Investigation
Methodology
Experiments
Issues For The . . .
Conclusion
Title Page
Page 12 of 29
Go Back
Full Screen
Close
Quit

5.2. Ideal Condition

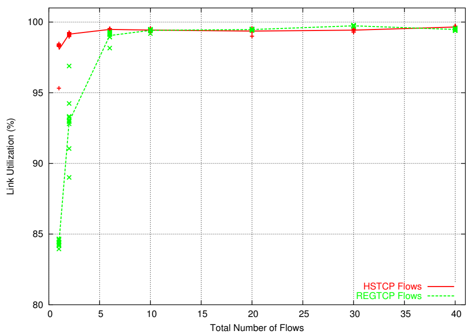

Figure 4: **Aggregated Link Utilization - Ideal Condition - Homogeneous Flows - RED**

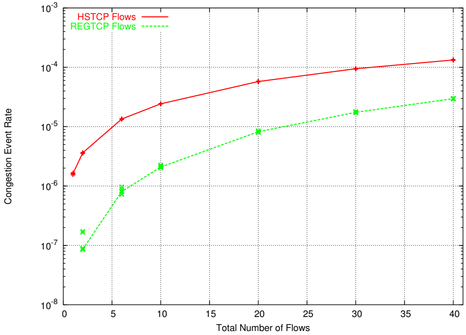

Figure 5: **Congestion Event Rate - Ideal Condition - Homogeneous Flows - RED**

- HSTCP needs less flows than REGTCP to reach full link utilization
- HSTCP produces a higher congestion event rate
- congestion event rate for HSTCP is never lower that 10^{-6}

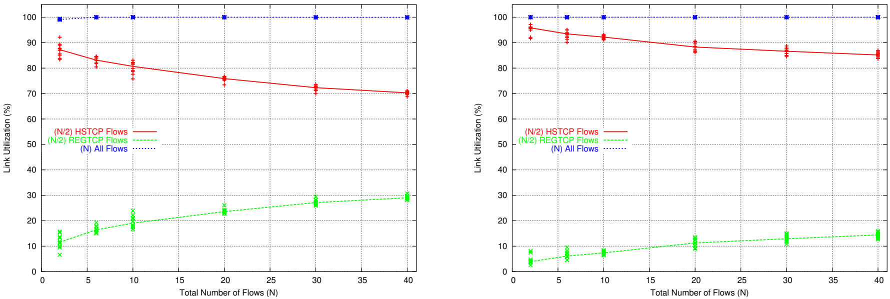

(a) **RED** (b) **DT**

Figure 6: **Aggregated Link Utilization - Ideal Condition - Heterogeneous Flows**

- the performance is separated by flow type
- bandwidth share used by HSTCP is higher than the bandwidth used by the REGTCP flows, independent of the type of router queue management used
- bandwidth share used by HSTCP decreases as the total number of flows increase

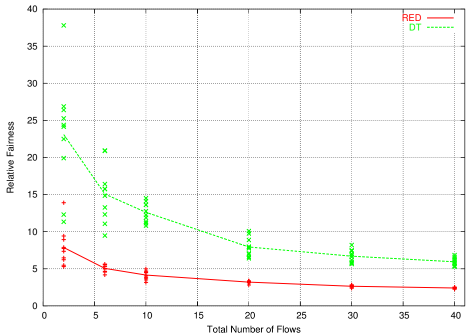

Figure 7: **Aggregated Relative Fairness - Ideal Condition**

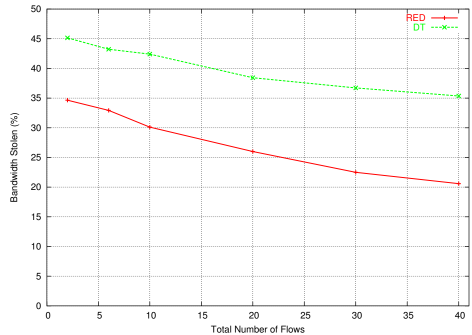

Figure 8: **Aggregated Bandwidth Stolen - Ideal Condition**

- fairness improves as the number of flows increases
- the relative fairness is lower when RED is deployed than when DT is used
- bandwidth stolen decreases as the number of flows increases
- HSTCP aggressiveness adapts as the traffic condition changes
- the distance between the amounts stolen between RED and DT increases slightly

Title Page
Page 15 of 29
Go Back
Full Screen
Close
Quit

5.3. Lossy Link Condition

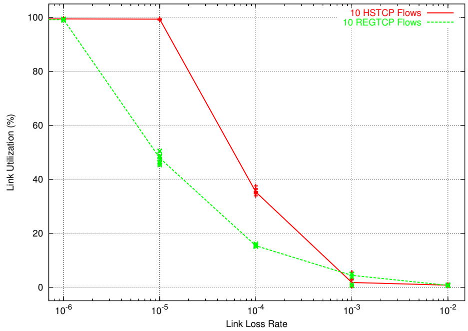

Figure 9: **Aggregated Link Utilization - Lossy Link Condition - Homogeneous Flows - RED**

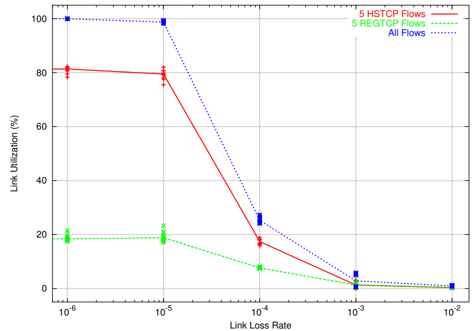

Figure 10: **Aggregated Link Utilization - Lossy Link Condition - Heterogeneous Flows - RED**

- The set of REGTCP flows show a rapid performance loss as the link loss rate increases
- the HSTCP flows showed better resilience to moderate link loss, and consistently used more bandwidth than the REGTCP flows
- for the set of mixed flows, the difference between the bandwidth used by the HSTCP flows and that used by the REGTCP flows decreases as the number of losses increases
- for a link loss rate around 10^{-5}, the link is fully utilized, and below this rate, congestive losses will be dominant

5.4. Bursty Traffic Condition

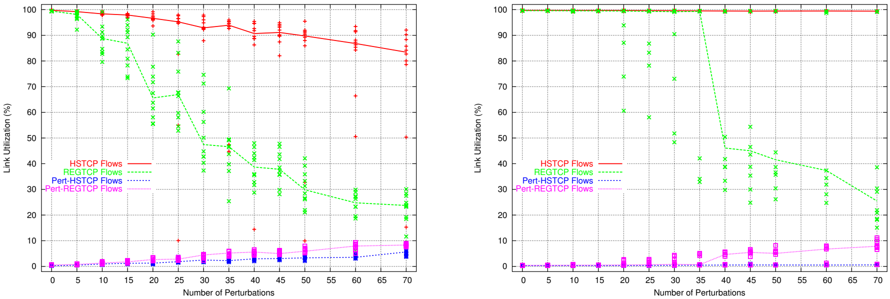

(a) **RED** (b) **DT**

Figure 11: **Aggregated Link Utilization - Bursty Traffic Condition - Homogeneous Flows**

- in RED case, the set of HSTCP flows decreases their link utilization smoothly and slowly as the number of perturbations increase
- the impact on the set of REGTCP flows is higher, and their performance goes down quickly as the number of perturbations increases
- strong impact of the use of distinct router queuing management
- the link utilization of the set of HSTCP flows is almost immune to the perturbations when the DT router queuing policy is used

Summary
What is it?
Topics of Investigation
Methodology
Experiments
Issues For The . . .
Conclusion
Title Page
Page 17 of 29
Go Back
Full Screen
Close
Quit

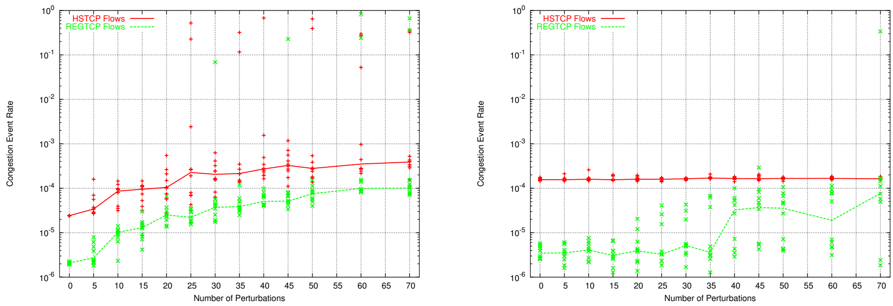

(a) **RED** (b) **DT**

Figure 12: **Congestion Event Rate - Bursty Traffic Condition - Homogeneous Flows**

- the congestion event rate increases continuously as the number of perturbations increases, when RED router queue management is used
- when DT router queue management is deployed, the set of HSTCP flows presents an almost constant congestion event rate
- the set of REGTCP flows has two levels of congestion event rates, probably caused by the occurrence of global synchronization

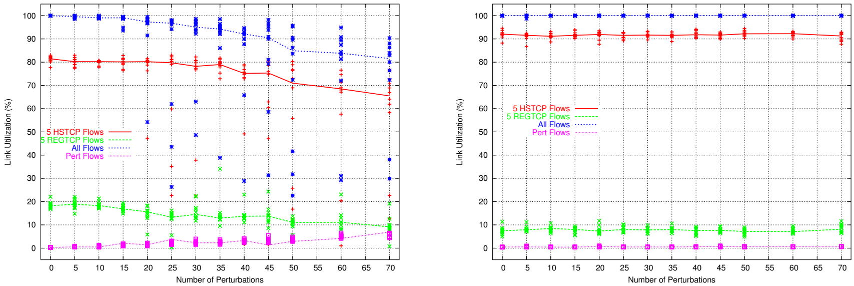

(a) RED (b) DT

Figure 13: **Aggregated Link Utilization - Bursty Traffic Condition - Heterogeneous Flows**

- poor link utilization of the set of REGTCP flows, however, this performance remains relatively constant as the number of perturbations increases

Title Page
Page 19 of 29
Go Back
Full Screen
Close
Quit

5.5. Competition among Heterogeneous Flows

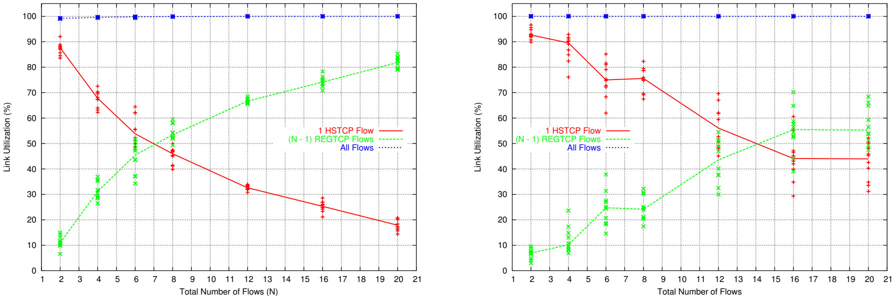

(a) **RED** (b) **DT**

Figure 14: **Aggregated Link Utilization - Competition Among Heterogeneous Flows**

- the HSTCP flow adapts with the amount of REGTCP flows used, and it avoids allowing the link to become idle
- the crosspoint of the HSTCP line and the REGTCP line shows the number of REGTCP flows that have equivalent performance to 1 HSTCP flow
- this number appears to be highly dependent on the type of router queue management used

5.6. Constant Link Loss Rate of 10^{-5}

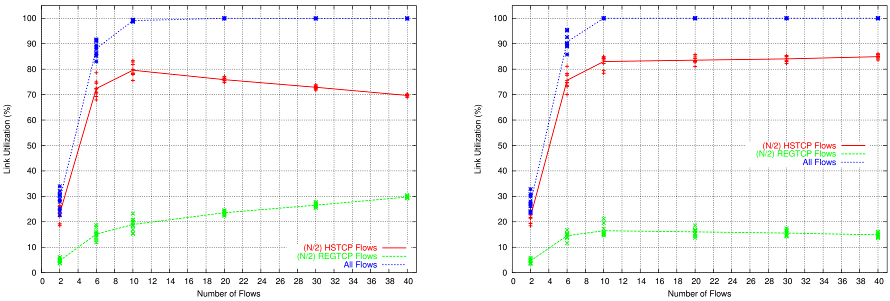

(a) **RED** (b) **DT**

Figure 15: **Aggregated Link Utilization - Constant Link Loss Rate of 10^{-5} - Heterogeneous Flows**

- the router queue management has influence on the link utilization for each type of flow
- while for RED, the link utilization for the HSTCP flows decreases as the total number of flows increases, for DT, the same link utilization stays constant or even slightly increases, as the number of flows increase

5.7. Parallel Streams on Lossy Link Condition

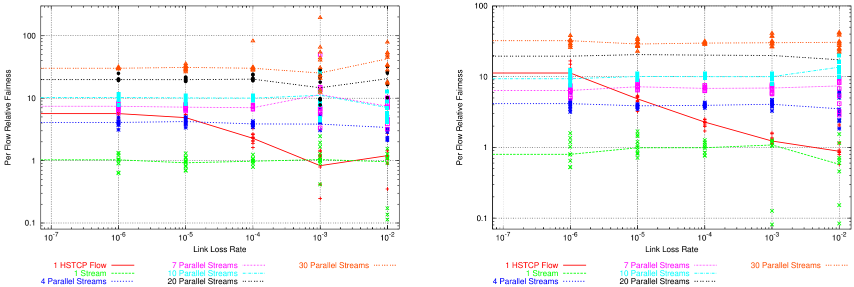

(a) RED (b) DT

Figure 16: **Per Flow Relative Fairness - Parallel Streams on Lossy Link Condition**

- when parallel streams are deployed, the relative fairness is almost constant over a wide range of link loss rates
- the relative fairness when HSTCP is used is not constant and has a wide range of values
- the variation of relative fairness of HSTCP is wider with DT than RED

5.8. Parallel Streams on Bursty Traffic Condition

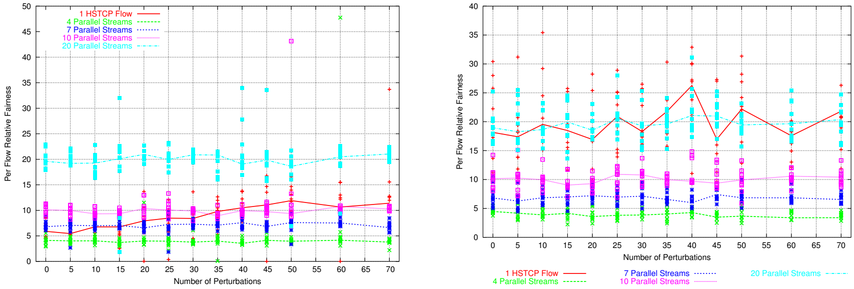

(a) RED (b) DT

Figure 17: **Per Flow Relative Fairness - Parallel Streams on Bursty Traffic Condition**

- when RED is used the relative fairness increases as the number of perturbations increase
- in DT case the relative fairness spreads over a wide range of values, but the average seems to be constant

Summary
What is it?
Topics of Investigation
Methodology
Experiments
Issues For The . . .
Conclusion
Title Page
Page 23 of 29
Go Back
Full Screen
Close
Quit

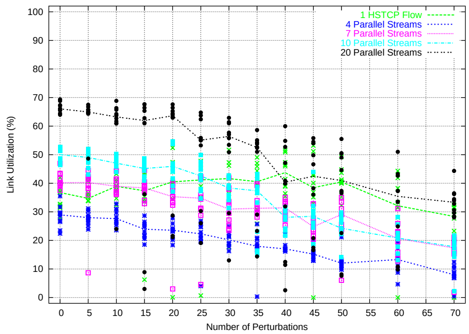

Figure 18: **Aggregated Link Utilization of Competing Parallel Streams - Parallel Streams on Bursty Traffic Condition - RED**

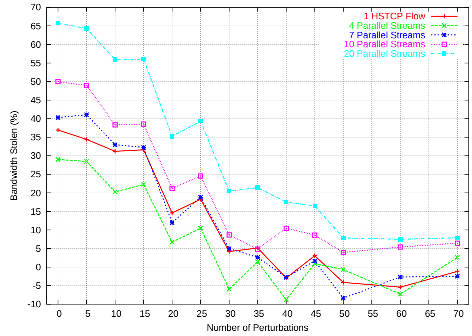

Figure 19: **Aggregated Bandwidth Stolen - Parallel Streams on Bursty Traffic Condition - RED**

- the HSTCP flow improved its performance as the number perturbation increased (RED case) until around 40 perturbations
- this improvement didn't happen stealing bandwidth from the 10 long-lived flows

6. Issues For The Deployment of HighSpeed TCP

6.1. Comparison with Regular TCP

- HighSpeed TCP performs better than Regular TCP for high-speed long-distance links
- In the presence of systemic losses, Regular TCP flows show poor link utilization
- a link loss rate between 10^{-5} and 10^{-4} prevented Regular TCP from making reasonable use of the link bandwidth available (less than 50% in our case)
- In this range, the HighSpeed TCP flows were able to use almost double the bandwidth used by Regular TCP

6.2. Fairness Impact

- The bandwidth share used by the HighSpeed TCP flows was higher than that used by Regular TCP flows when both types of flows competed for the same link, but the amount of the link bandwidth used decreases as the total number of flows increased
- This adaptability is very interesting in the context of high speed links. It avoids having a link become idle caused by the slow dynamic of Regular TCP, and yet it does not prevent more Regular TCP streams from obtaining a reasonable share of the link
- Bursty traffic had only a small influence in the amount of bandwidth that the HighSpeed TCP flows stole from the Regular TCP flows, thus it had little influence on the fairness

6.3. Effects of Router Queue Management

- the change of the queue management scheme did not significantly affect the link utilization of HighSpeed TCP flows in most cases
- the difference on router queue management policies was clear when HighSpeed TCP was submitted to bursty traffic. The link utilization for HighSpeed TCP flows decreased slightly with RED, but it was not affected when DT router queue policy was used
- the difference from RED to DT was the higher amount of bandwidth that HighSpeed TCP flows took from Regular TCP flows when DT was used

6.4. Use For Bulk Data Transfer

- Deployment of HighSpeed TCP requires changes to the TCP stack, but once these changes are made all applications can benefit from them
- For parallel streams it is necessary to change the application programs and to know a priori the number of parallel flows to transmit
- HighSpeed TCP presents better fairness and adaptability to an environment of variable congestion event rates than parallel streams

7. Conclusion

- HighSpeed TCP overforms TCP, and has an adaptability that makes easy an incremental adaption approach
- HighSpeed TCP is easy to deploy avoiding changes in routers and programs
- HighSpeed TCP is appropriate to bulk data transfer applications, because it is able to keep high throughput in different network conditions